arXiv:1407.4760v1 [math.OC] 17 Jul 2014

What Makes a Good Plan? An Efficient Planning Approach to Control Diffusion Processes in Networks

Kevin Scaman Argyris Kalogeratos Nicolas Vayatis
CMLA – ENS Cachan, CNRS, France
{scaman, kalogeratos, vayatis}@cmla.ens-cachan.fr

July 18, 2014

Abstract

In this paper, we analyze the quality of a large class of simple dynamic resource allocation (DRA) strategies which we name *priority planning*. Their aim is to control an undesired diffusion process by distributing resources to the contagious nodes of the network according to a predefined *priority-order*. In our analysis, we reduce the DRA problem to the *linear arrangement* of the nodes of the network. Under this perspective, we shed light on the role of a fundamental characteristic of this arrangement, the *maximum cutwidth*, for assessing the quality of any priority planning strategy. Our theoretical analysis validates the role of the maximum cutwidth by deriving bounds for the extinction time of the diffusion process. Finally, using the results of our analysis, we propose a novel and efficient DRA strategy, called *Maximum Cutwidth Minimization*, that outperforms other competing strategies in our simulations.

1 Introduction

Diffusion processes usually arise in systems involving agents whose behaviors depend on their close environments. Diseases, computer viruses, ideas and interests are spread through a social network by means of interactions among its users. In all these phenomena, a change in one agent may affect the actions of other neighboring agents, resulting, under certain conditions, in a massive change of behavior at the network scale.

A fair body of articles consider the problem of influence maximization, which attempts to maximize the spread of a diffusion process. On the other hand, being able to suppress or remove an undesired information or social diffusion process has received less attention, even though critical in many real-life situations. In public health, scenarios in which the spread of a virus needs to be controlled have been extensively studied by epidemiologists. Various analogues emerged in

modern information networks in which a diffusion can be engineered to be viral and may cause huge negative social and economic effects. For instance, during the London riots during the summer of 2011, public opinion was obfuscated and baffled by many false rumors [1]. During large scale natural disasters, there is also a risk to follow misinformation diffused through social media by various individuals, while trying to coordinate some rescue and volunteer teams [2, 3]. Another recent example arose in March 2013, when a false tweet from the account of Associated Press was headlined with "*Breaking: Two explosions in the White House and Barack Obama is injured*". The false rumor was immediately retweeted by about 2 million users causing panic and speculation for a few minutes in the US stock market. As a matter of fact, negative publicity is highly damaging for organizations and brands: even though ill-founded, a slander can dramatically affect the interests of a company due to the massive scale of the buzz-like effect.

The control of diffusion processes has been studied in various fields, including epidemiology and computer networks resilience; the respective literature can generally be divided in three complementary lines of research:

- *Static vaccination strategies.* Most of the epidemic literature focuses on static control actions such as permanently removing a set of edges or nodes of the network [4, 5, 6, 7, 8]. In this case, the available budget is considered fixed, and the effect of a control action *permanent*. Examples of static resource allocation strategies can be found in [9, 10].

- *Budget optimization.* Complementary to resource allocation, the determination of the right budget of resources to be spent each time step, in order to fulfill cost and efficiency constraints, has critical impact on the resulting strategy. Several studies lay on the line of optimal budget estimation, assuming that the network administrator is capable of storing resources for later use [11, 12, 13]. These approaches usually make the simplifying assumption of *uniform mixing*, i.e. that the contagious nodes are uniformly scattered in the network. Therefore, they do not address the problem of how exactly to allocate the resources on the nodes of the network, but rather how many resources should be provided at each time to cause a desired macroscopic result.

- *Dynamic resource allocation.* A few studies consider dynamic strategies for allocating resources for dealing with epidemics. One of the most well-known such strategy is *contact-tracing* [14] that consists in healing the neighbors of contagious nodes. In practice, this approach has been shown inefficient to contain epidemics [14], especially when they are beyond a very initial state.

Our major contribution in this article is to introduce a particular class of strategies for suppressing an undesired diffusion. Instead of choosing a set of nodes whose behavior will be permanently modified, we allow the network administrator to change the distribution of the *resources* during the diffusion. In

other words, we consider *targeted* and *temporal* action on individual nodes of the network, that can affect their behavior. Since reacting to fast spreading phenomena is difficult to achieve, we consider a simple set of dynamic strategies relying on a predefined *priority-order*. The strategy will gradually suppress the diffusion and finally remove the undesired contagion by focusing on the first contagious users according to the priority order.

In what follows, Sec. 2 describes the model used for the diffusion process as well as the control actions available to the network administrator. Sec. 2.2 presents the idea of *priority planning* as well as a natural representation of the problem connecting our analysis to *linear arrangement* problems. This anaylysis sheds light on the role of the network's *maximum cutwidth* for the efficiency of a given strategy. By minimizing this value, we develop an efficient strategy in Sec. 3, and validate in Sec. 4 that our intuition is valid by deriving theoretical results on the extinction time of the diffusion process. Finally, we present experimental results in Sec. 5 and show that: i) The derived bounds are very close to the epidemic threshold, thus validating the fundamental role of the *maximum cutwidth* in the evaluation of such strategies. ii) The proposed strategy outperforms its competitors in a wide range of scenarios.

2 Diffusion and control model

2.1 The *Susceptible-Infected-Susceptible* epidemic model

We will consider a simple diffusion process known in the epidemiology literature as *N-Intertwined Susceptible-Infected-Susceptible* (SIS) model [15]. According to this model, a diffusion can be spread through the edges of the network and turn the state of nodes from *susceptible* (or *healthy*) to *infected*. When a node becomes infected, it can in turn spread the contagion to its direct neighbors and, after some amount of time, return to the susceptible state without bearing any immunity. At each time, the network administrator can take control actions in order to reduce the epidemic. These actions are represented as a *budget* of *resources* (or treatments in the epidemic analog) to distribute in the network. Each resource increases the recovery rate of the receiver node. In a continuous-time framework, the distribution of control resources can be revised anytime. However, in practice the situations in which this is needed is only when there is a change in the state of the network, i.e. a new node infection or recovery, or a modification in the available resource budget.

This model is well suited to situations in which an undesired contagion affects a network, and the control action is *local* and *expensive*. Among other application examples, controlling epidemics using antidotes, limiting rumors via targeted action or allocating resources geographically to fight against a societal problem, seem valid scenarios for such a diffusion and control model.

Formal definition. Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be a network of N nodes with adjacency matrix A, where $A_{ij} = 1$ if $i \neq j$ and edge $(i, j) \in \mathcal{E}$, else $A_{ij} = 0$. We describe the state of the diffusion process with a *state vector* $X(t) \in \mathbb{R}^N$ that keeps the

state of each node of the network: $X_i(t) = 1$ if node i is contagious at time t, else $X_i(t) = 0$. We also describe the control action on the network with a *resource vector* $R(t)$, where $R_i(t)$ is 1 iff node i is given a resource at time t. In such case, and following the epidemic analogy, we say that node i is being *healed* by the resource. Using also the formalism of [16], we model the diffusion with a *continuous-time Markov process* which has the transition rates:

$$\begin{array}{l} X_i(t) : 0 \to 1 \text{ at rate } \beta \sum_j A_{ji} X_j(t); \\ X_i(t) : 1 \to 0 \text{ at rate } \delta + \rho R_i(t), \end{array} \tag{1}$$

where β, δ, ρ are, respectively, the transmission rate over one network edge, the recovery rate without receiving a resource, and the increase in recovery rate that a resource induces. Note that the diffusion process is continuous in time and thus $X(t)$ is a stochastic process. The action of the resource on a node increases its chances to return to the *healthy* state. Finally, we define two dimensionless parameters: $r = \frac{\beta}{\delta}$ the *effective spreading rate* of the diffusion, and $e = \frac{\rho}{\delta}$ the *resource efficiency*.

A *control strategy* takes as input the network $\mathcal{G}$, the characteristics of the diffusion (β, δ, and ρ) and the network state $X(t)$, and returns the distribution of the resources $R(t)$ (Eq. 1). In the following, we refer to the problem of finding an optimal control strategy with respect to the minimization of the spread of the diffusion process as the *dynamic resource allocation* (DRA) problem. We consider as *budget* $b(t)$ the maximum number of resources that can be distributed in the network at time t, and that the available budget at time t cannot be stored for later use. Following the epidemiology literature, we will denote as *epidemic threshold under a given strategy* the resource efficiency e above which the control actions removes the contagion in *reasonable time*, that is less than exponential in the number of nodes of the network.

2.2 Priority planning: a control plan to gradually remove a contagion

2.2.1 Definition

A *priority planning* is a DRA strategy that distributes resources to the top-$b(t)$ infected nodes according to a fixed *priority ordering* of the nodes, where $b(t)$ is the available budget of resources at each time. In mathematical terms, an *ℓ-priority planning* is defined by a bijective mapping $\ell : \mathcal{V} \to \{1, ..., N\}$ of the N nodes of the network s.t. $\ell(v)$ is the position of node v in the priority order. The strategy then selects the first $b(t)$ infected nodes according to ℓ:

$$R_i(t) = \begin{cases} 1 & \text{if } X_i(t) = 1 \text{ and } \ell(v_i) \leq \theta(t); \\ 0 & \text{otherwise,} \end{cases} \tag{2}$$

where $\theta(t)$ is an adjusted threshold set so that $\sum_i R_i(t) = b(t)$.

These strategies may be regarded as simple planning strategies for the removal of an undesired contagion: the order in which the nodes will be healed is

determined prior to the beginning of the diffusion. Then, this order is respected no matter how the diffusion process evolves, and the strategy removes the contagious nodes gradually starting from the first in the priority-order to the last. Although our strategy (Sec. 3) can handle a generic time-dependent budget, we only consider a constant budget rate $b(t) = b_{tot}$ for theoretical developments as well as focused experimental evaluation.

2.2.2 Interpretation of priority planning

In this section, we present a novel perspective on analyzing the DRA problem which leads to efficient priority planning strategies. More specifically, we reduce the problem of determining a good priority-order to a *linear arrangement* (LA) problem.

Linear arrangement. Formally, a linear arrangement $\ell \colon \mathcal{V} \to \{1, ..., N\}$ maps the nodes of $\mathcal{G}$ on N discrete positions located on a line by assigning one position-label to each node (Fig. 1). LA is a class of combinatorial optimization problems whose purpose is to minimize some functional ϕ over the space $\mathcal{L}$ of all possible node permutations: $\ell^* = \text{argmax}_{\ell \in \mathcal{L}}\, \phi(\mathcal{G}, \ell)$. The problems in the LA class are also referred to as *graph layout* problems (for a survey see [17]), and indicative applications are the graph drawing, VLSI design, and network scheduling.

Probably the most popular LA instance is the *minimum p-sum linear arrangement* (MpLA, or MLA when $p = 1$) [18] that minimizes the following functional:

$$\text{M}p\text{LA}\colon \quad \phi_1(\mathcal{G}, \ell) = \Big(\textstyle\sum_{i,j} A_{ij} \, |\ell(v_i) - \ell(v_j)|^p \Big)^{1/p}. \tag{3}$$

In other words, for $p = 1$, this minimizes the weighted sum over all distances between node assignements $\ell(u)$ to the N discrete equally spaced positions. Another category of LA problems is related to the minimization of some maximum value of the LA, for example the *bandwidth* (i.e. the maximum edge length), or the *workbound* (i.e. the sum of maximum edge length of each node). There is also the *minimum cutwidth linear arrangement* (MCLA) problem, where the *cutwidth* at a location c lays among the node positions c and $c+1$ of the LA and is the weighted edge-cut at that position. Then, the minimized functional is the maximum of the $N-1$ local cuts, and more precisely:

$$\text{MCLA}\colon \quad \phi_2(\mathcal{G}, \ell) = \max_{c=1,...,N-1} \textstyle\sum_{i,j} A_{ij} \mathbb{1}_{\{\ell(v_i) \le c < \ell(v_j)\}}, \tag{4}$$

where $\mathbb{1}_{\{\cdot\}}$ is the indicator function that returns one if the input condition is true, else zero. For the simplicity of notations we also denote the value of the above functional as $\mathcal{C}_{max} = \phi_2(\mathcal{G}, \ell)$. Note that symmetric LAs are evaluated as of equal quality by ϕ_1, ϕ_2. Fig. 1 illustrates two LAs of a small network and an example of the cost values derived in each case by Eq. 3 and Eq. 4. Among them, the arrangement of Fig. 1b is better, indeed optimal, in terms of both objective functions ϕ_1, ϕ_2.

Designing an efficient control plan. The analysis of the diffusion process under the perspective of a linear arrangement, sheds light on how to design a

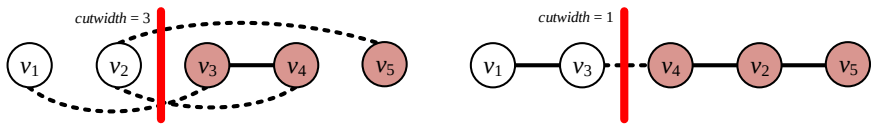

(a) LA with mapping $\ell\colon \mathcal{V} \to \{1,2,3,4,5\}$ (b) LA with mapping $\ell'\colon \mathcal{V} \to \{1,3,4,2,5\}$

Figure 1: The priority orders of two static strategies are visualized as the linear arrangements of the nodes based on two mappings ℓ and ℓ'. Healthy nodes (white) and contagious nodes (red) form two distinct subgraphs that are connected by the *contagious edges* (dashed lines) of the *front* illustrated as a red vertical line. (a) M*p*LA cost, for $p = 1$, $\phi_1(\mathcal{G},\ell) = 8$ and MCLA cost $\phi_2(\mathcal{G},\ell) = 3$; (b) $\phi_1(\mathcal{G},\ell') = 4$ and $\phi_2(\mathcal{G},\ell') = 1$, respectively.

DRA control plan to suppress a diffusion efficiently. More specifically, knowing that the contagion can only be spread through the *contagious edges* which connect infected with healthy nodes, and that any DRA plan is based on a static priority-order for healing the nodes, we realize that it is highly critical to have as less contagious edges as possible during the whole diffusion process.

Fig. 1 represents two diffusion instances that are being suppressed according to different plans; we consider that each priority-order is deployed from the left to the right. The current state of the plan is denoted with a red vertical line, the *front*, that roughly separates the healthy from the contagious nodes and indicates where the strategy would allocate resources. In practice, healthy nodes can also appear on the right side of the front due to self-recovery. The number of edges crossing through the front, i.e. the cutwidth at that position, indicates how vulnerable are actually the healthy nodes of the cleared part of the network laying left to the front. Therefore, we can argue that a good solution to the M*p*LA minimization would provide a smoother plan that is generally *easier to proceed from state-to-state*, while the MCLA would minimize the $\mathcal{C}_{max}$ which is *the most difficult state of the plan* and, thus, a determinant for its eventual success in completely removing the contagion. This implies that the minimization of the $\mathcal{C}_{max}$ can be a proxy for minimizing the epidemic threshold for any arbitrary network. Next, we theoretically show that these intuitive remarks are valid.

3 The *Maximum Cutwidth Minimization* control strategy

Based on the analysis of the previous section, that uncovered a strong dependency between the epidemic threshold of the diffusion process and the *maximum cutwidth*, we propose a new DRA algorithm for arbitrary networks where the main idea is to distribute resources to contagious nodes in the order that minimizes $\mathcal{C}_{max}$. Specifically, given a network $\mathcal{G}$, we compute, prior to the diffusion process, the linear arrangement $\ell_{MC}(\mathcal{G})$ of its nodes with minimum $\mathcal{C}_{max}$ value

Algorithm 1 *Maximum cutwidth minimization* (MCM) control strategy

Input : network $\mathcal{G}$, infection state vector $X(t)$, budget size $b(t)$
Output: the resource allocation vector $R(t)$

▷ **Prior to the diffusion:**
Compute the priority-order $\ell = \ell_{MC}(\mathcal{G})$ by minimizing the maximum cutwidth $\mathcal{C}_{max}$

Order the nodes of $\mathcal{G}$ according to ℓ, i.e. compute the node list $(v_1, ..., v_N)$ s.t. $\forall i \in \{1, ..., N\}$, $\ell(v_i) = i$

▷ **During the diffusion:**
if $\sum_i X_i(t) < b(t)$ **then**
  return $X(t)$
end if
Let $R(t) = \mathbb{0}$ a zero N-dimensional vector, also $i = 1$ and $budget = b(t)$
while $budget > 0$ **do**
  if $X_{v_i}(t) = 1$ **then**
    $R_{v_i}(t) \leftarrow 1$
    $budget \leftarrow budget - 1$
  end if
  $i \leftarrow i + 1$
end while
return $R(t)$

using any available optimization algorithm for this problem. Then, during the diffusion, the strategy distributes the budget of resources to the contagious nodes according to the order of $\ell_{MC}(\mathcal{G})$. The pseudocode of the strategy is summarized in Alg. 1, while theoretical justification and results are provided in Sec. 4.

Maximum cutwidth linear arrangement. MCLA is known to be an NP-hard problem. However, approximation heuristics do exist in the related literature [19, 20]. One of the difficulties of this problem is that the cost function to optimize (Eq. 4) is extremely flat in the search space, due to the fact that slight changes in the arrangement will most probably not change $\mathcal{C}_{max}$. For this reason, we chose to relax the MCLA problem by optimizing the *p-sum linear arrangement* problem with $p = 1$ (MLA, see Sec. 2.2.2), which is easier than MCLA and more suited to gradient descent or simulated annealing methods. Furthermore, as discussed in Sec. 2.2.2, MLA may produce a smooth priority-order that exhibits some desirable properties.

Scalability of the method. The scalability of the strategy is highly dependent on the chosen algorithm for finding the optimal order. For our experimental results, we applied a hierarchical approach to take advantage of the group-structure of the social network: i) first, we identified dense clusters by applying *spectral clustering* and we ordered the clusters (considered as high-level nodes) using *spectral sequencing* [21], ii) then, we computed a good ordering of the nodes in each of the clusters independently using spectral sequencing fol-

lowed by an iterative approach based on random node swaps, inspired by [22], iii) finally, the swap-based approach was reapplied to optimize the ordering all together. The whole process achieved fairly good results (see Tab. 1) in reasonable time. Since clustering and spectral sequencing depend on the computation of eigenvectors for the highest eigenvalues of an $N \times N$ sparse matrix with $|\mathcal{E}|$ non-zero entries, the overall algorithm has a complexity $O(|\mathcal{V}| + |\mathcal{E}|)$ [23]. Hence, MCM is generally scalable to the size of real social networks.

4 Theoretical bounds for the extinction time and the epidemic threshold

We now provide a justification of the design of our MCM algorithm. The following theorem gives an *upper bound* for the expected extinction time under any priority planning in the simple case of $b(t) = 1$. The general case, while much more complex to analyze, is very similar in our experiments provided that the treatment efficiency e is multiplied by the available budget $b(t)$ (see Sec. 5). Above a threshold value that depends on the *maximum cutwidth* $\mathcal{C}_{max}$ of the network under a considered priority-order, Theorem 1 proves that the diffusion process converges in reasonable time to its absorbent state.

Theorem 1. *Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be a network of N nodes and assume a fixed budget $b(t) = 1$. Let $\ell : \mathcal{V} \to \{1, ..., N\}$ be an ordering of the nodes of $\mathcal{G}$, and $\mathcal{C}_{max}$ be the maximum cutwidth of ℓ, or equivalently the highest number of contagious edges during the planned removal of the contagion. Then the following upper bound holds for the expected extinction time $\mathbb{E}[\tau]$ under the ℓ-priority plan and starting from a total infection:*

If $\rho > \beta \mathcal{C}_{max} (1 + 2\sqrt{\epsilon} + \epsilon) - \delta$, then

$$\mathbb{E}[\tau] \;\leq\; \frac{N}{\rho + \delta - \beta \mathcal{C}_{max} (1 + 2\sqrt{\epsilon} + \epsilon)} \tag{5}$$

where $\epsilon = \frac{d_{max}(1+\ln N)}{\mathcal{C}_{max}}$ and d_{max} is the maximum node degree of the network.

This bound relates the extinction time to the number of contagious edges in the worst step of the strategy (i.e. its *maximum cutwidth*). When $\mathcal{C}_{max}$ is such that $d_{max}(1 + \ln N) \ll \mathcal{C}_{max}$, this formula bounds the epidemic threshold under a given priority planning by $r\mathcal{C}_{max} - 1$, where $r = \beta/\delta$ is the *effective spreading rate*, thus giving a special emphasis on the role of $\mathcal{C}_{max}$ for assessing the quality of a given strategy.

Corollary 1. *Under the same hypotheses as Theorem 1, the epidemic threshold e^*, i.e. the lowest resource efficiency $e = \frac{\rho}{\delta}$ s.t. the diffusion converges to zero in reasonable time, is upper bounded:*

$$e^* \;\leq\; r\mathcal{C}_{max} \left(1 + 2\sqrt{\epsilon} + \epsilon\right) - 1, \tag{6}$$

where $\epsilon = \frac{d_{max}(1+\ln N)}{\mathcal{C}_{max}}$ and $r = \frac{\beta}{\delta}$ is the effective spreading rate.

While ϵ is necessary for the theorem to hold, we rather consider it as an artifact of the proof and not a fundamental aspect of the result. As such, we expect $r\mathcal{C}_{max} - 1$ to be closer to the epidemic threshold (see specifically Sec. 5.1 for an experimental validation of this intuition). This result verifies that, according to the intuition provided in Sec. 2.2.2, removing an undesired contagion requires the resource efficiency to be *as high as needed for the worst step* of the specified plan. Equivalent to the celebrated relationship between the epidemic threshold under no control strategy and the spectral radius of the adjacency matrix, this result is fundamental for understanding the behavior of the diffusion process and designing efficient DRA strategies. The simulations in Sec. 5, attest that minimizing this upper bound is very efficient for dynamically controlling a diffusion process.

5 Experimental results

5.1 Quality of the theoretical bound

In Fig. 2, we show that the relationship between the *maximum cutwidth* $\mathcal{C}_{max}$ and the epidemic threshold under a specific priority plan is very stable, nearly linear, and hence $\frac{r\mathcal{C}_{max}}{b_{tot}}$ is in fact a very good approximation. Each plotted point is the epidemic threshold under a given strategy and for a given network instance. We sampled 100 networks of 1000 nodes from 5 random network generators (e.g. see [24] for details on these network types): i) Erdös-Rényi , ii) Preferential attachment, iii) Small-world, iv) Geometric random [25] and v) 2D regular grids. We also sampled DRA control strategies at random among the set of strategies described next in Sec. 5.2, in order to cover a wide range of different scenarios.

Fig. 2 is a summary of these results. The epidemic threshold is always below $\frac{r\mathcal{C}_{max}}{b_{tot}}$, but it stays very close to this value and, more importantly, shows a close to deterministic behavior with respect to this value even in the case of low infectivity (small r value) where intuitively the random self-recovery of nodes become more significant. This result justifies the minimization of $\mathcal{C}_{max}$ as a proxy for minimizing the epidemic threshold, and thus remove the contagion more efficiently.

5.2 Application on a real social network

For evaluating our strategy, the *maximum cutwidth minimization* (MCM), we considered a realistic scenario where an SIS epidemic is being spread in a subset of the Twitter social network, denoted as TwitterNet [1]. TwitterNet consists of $1,000$ ego-networks extracted from the social network [26]. In order to perform experiments in the setting considered in this article, we symmetrized this network and obtained an undirected network of $81,306$ nodes and $1,342,303$

[1] Available on: http://snap.stanford.edu/data/index.html.

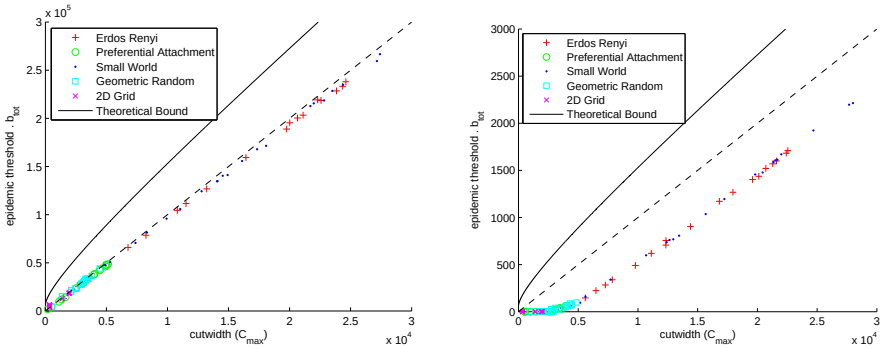

(a) high infectivity scenario: $r = 10$ (b) low infectivity scenario: $r = 0.1$

Figure 2: Epidemic threshold w.r.t. the maximum cutwidth $\mathcal{C}_{max}$ for various random network types of $N = 1000$ nodes. For the theoretical bound, $d_{max} = 100$ was used. The dashed line indicator has slope equal to the effective spreading rate r.

Strategy	**Max cutwidth**	**Max cutwidth % MCM**	**Expected epidemic threshold** ($r = 0.1$, $b_{tot} = 100$)
RAND	670,000 ± 1000	931%	670
MN	628,571	874%	629
LN	628,571	874%	629
LRSR	349,440	486%	349
MCM	**71,956**	**100%**	**72**

Table 1: Maximum cutwidth values ($\mathcal{C}_{max}$) for different strategies in the TwitterNet used in our experiments. The expected epidemic threshold is $\frac{r\mathcal{C}_{max}}{b_{tot}}$ based on Corollary 1 and the experiments of Sec. 5.1.

edges. The resulting network has a single connected component and contains a rich community structure.

The diffusion control literature is mainly focused on two types of strategies: static vaccination strategies and dynamic strategies based on a uniform mixing hypothesis. We compare MCM to the first type of strategies by considering state-of-the-art vaccination algorithms and use their output as a priority-order of the nodes, and to the second by considering random allocation of resources. As discussed in Sec. 2, the problem of finding the optimal budget through time is complementary to our approach. From this perspective, one of the primary question we address here is whether targeting specific nodes in the network can lead to substantial gains in the efficiency of the method, compared to the random distribution of the resources to the contagious nodes. Fig. 3a and 3b compare the MCM strategy against the following four baseline approaches for creating a priority-order:

- *Random order* (RAND): the order is a random permutation of the nodes of the network.
- *Most neighbors* (MN): this strategy gives priority to high degree nodes.

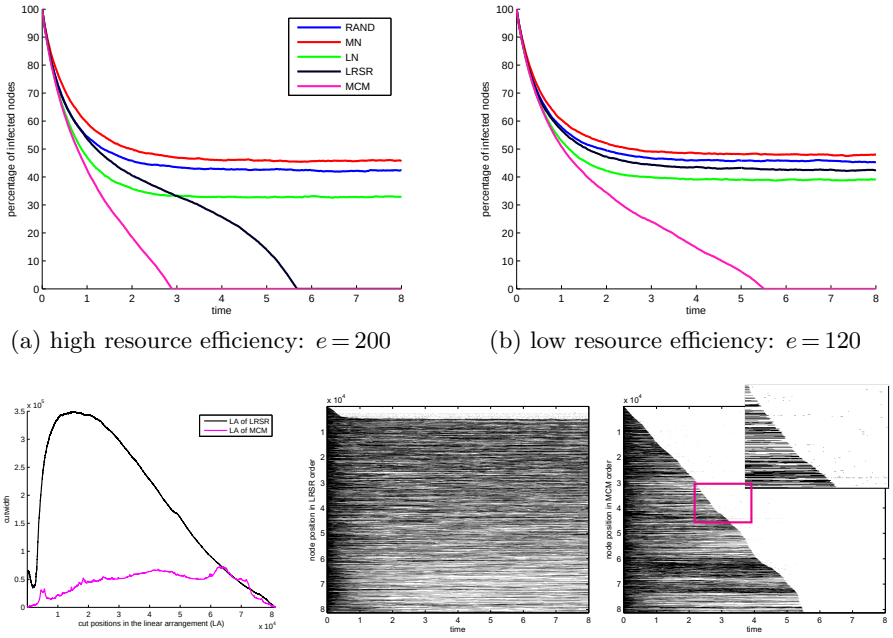

(a) high resource efficiency: $e = 200$ (b) low resource efficiency: $e = 120$

(c) cutwidth values (d) network state – LRSR (e) network state – MCM

Figure 3: (a-b) Simulation of an undesired diffusion process in the TwitterNet subset of 81,306 nodes [26]. $\delta = 1$, $r = 0.1$ and $b_{tot} = 100$. MCM outperforms other heuristics. (c) Cutwidths for LRSR and MCM. (d-e) Visualization of the diffusion in (b) at the node level (contagious nodes in black). Nodes are ordered according to LRSR and MCM linear arrangements, respectively. A closer look of (e) is also provided in an inset figure.

Intuitively, it begins by removing the contagion from the core of the network, and gradually reaches its periphery.

- *Least neighbors* (LN): this strategy gives priority to low degree nodes. Conversely to MN, LN begins with the periphery of the network and converges to its central part.

- *Least reduction in spectral radius* (LRSR): this strategy is based on the vaccination literature and more specifically the work in [5]. This strategy gives priority to nodes whose removal will lead to the maximum decrease of the spectral radius of the adjacency matrix of the resulting network.

Concerning the applicability of Theorem 1, the maximum node degree in the network is $d_{max} = 3383$, which leads to a small ϵ value of 0.6 even for the MCM strategy (and $\epsilon = 0.1$ for LRSR). We can thus be relatively confident in approximating the epidemic threshold under a given priority planning by $\frac{r\mathcal{C}_{max}}{b_{tot}}$. Fig. 3c shows the cuts in the two best priority-orders, LRSR and MCM, while Tab. 1 summarizes the $\mathcal{C}_{max}$ values of the priority-orders produced by the

different compared strategies. Note that the $\mathcal{C}_{max}$ of the priority-order produced by MCM is *5 times smaller* than that of LRSR which is the best among its competitors. In essence, this implies that MCM would manage to suppress the diffusion process with resource efficiencies *5 times smaller* than LRSR.

The results in Fig. 3a and 3b illustrate that the proposed MCM strategy is more efficient than its competitors in removing the contagion from the network. In Fig. 3d and 3e, the evolution of the diffusion process is represented as follows: each line of the figure contains the state of one node of the network throughout the whole process (black for contagious and white for healthy). Nodes are sorted in y-axis according to the considered priority-order. We can observe that the cutwidth acts as a barrier for the LRSR: the large cutwidth values at the beginning of the LRSR order ($\frac{r\mathcal{C}_{max}}{b_{tot}} \approx e = 120$ for the 5000-th node of the linear arrangement) prevents the strategy from consistently removing the contagion from more than the first 5000 nodes of the plan. Note that healthy nodes also appear beyond the front, however this is not due to the control actions but rather due to self-recovery.

On the contrary, MCM gradually reduces the contagion and the advancement of the front is clearly visible. These results agree with our previous analysis, and show that: i) the uniform mixing hypothesis leads to a massive drop in efficiency, since MCM substantially outperforms the random allocation strategy, ii) while efficient in the static vaccination problem, centrality-based priority-orders are suboptimal for the DRA problem, iii) a good criterion for assessing the quality of a priority-order is actually in terms of its $\mathcal{C}_{max}$ value.

6 Conclusion

In this paper, we presented a novel type of dynamic strategies for allocating resources in a network, called *priority planning*, that aims to suppress an undesired contagion. We reduced the planning problem to that of linear arrangement of the nodes, and, based on theoretical analysis on the quality of any priority-order, i) we demonstrated the key role of the *maximum cutwidth* for assessing if a strategy would be eventually successful in removing the contagion, and ii) we derived a strategy, called *maximum cutwidth minimization* (MCM), that distributes resources to nodes in a priority-order with minimum maximum cutwidth. Our experimental results verified that, for a wide range of network types, the maximum cutwidth is indeed a good approximation of the epidemic threshold under a given strategy, and that the MCM strategy outperforms other competing strategies in a real-world social network.

Acknowledgments

This research is part of the SODATECH project funded by the French Government within the program of "*Investments for the Future – Big Data*".

References

[1] R. Procter, F. Vis, and A. Voss, "Reading the riots on twitter: methodological innovation for the analysis of big data," *Intern. J. Social Research Methodology*, vol. 16, no. 3, pp. 197–214, 2013.

[2] S. Vieweg, A. L. Hughes, K. Starbird, and L. Palen, "Microblogging during two natural hazards events: What twitter may contribute to situational awareness," in *Proc. of the SIGCHI Conf. on Human Factors in Computing Systems*, (New York, NY, USA), pp. 1079–1088, ACM, 2010.

[3] A. L. Hughes, L. A. A. St. Denis, L. Palen, and K. M. Anderson, "Online public communications by police & fire services during the 2012 hurricane sandy," in *Proc. of the 32nd ACM Conf. on Human Factors in Computing Systems*, (New York, NY, USA), pp. 1505–1514, ACM, 2014.

[4] R. Cohen, S. Havlin, and D. Ben-Avraham, "Efficient immunization strategies for computer networks and populations," *Physical Review Letters*, vol. 91, no. 24, p. 247901, 2003.

[5] H. Tong, B. A. Prakash, T. Eliassi-Rad, M. Faloutsos, and C. Faloutsos, "Gelling, and melting, large graphs by edge manipulation," in *Proc. of the 21st ACM Intl. Conf. on Information and Knowledge Management*, pp. 245–254, ACM, 2012.

[6] Y. Wang, D. Chakrabarti, C. Wang, and C. Faloutsos, "Epidemic spreading in real networks: An eigenvalue viewpoint," in *Proc. of the 22nd Intl. Sym. on Reliable Distributed Systems*, pp. 25–34, IEEE, 2003.

[7] C. M. Schneider, T. Mihaljev, S. Havlin, and H. J. Herrmann, "Suppressing epidemics with a limited amount of immunization units," *Physical Review E*, vol. 84, no. 6, p. 061911, 2011.

[8] V. M. Preciado, M. Zargham, C. Enyioha, A. Jadbabaie, and G. Pappas, "Optimal vaccine allocation to control epidemic outbreaks in arbitrary networks," *arXiv preprint arXiv:1303.3984*, 2013.

[9] V. M. Preciado, M. Zargham, C. Enyioha, A. Jadbabaie, and G. J. Pappas, "Optimal resource allocation for network protection: A geometric programming approach.," *CoRR*, vol. abs/1309.6270, 2013.

[10] F. Chung, P. Horn, and A. Tsiatas, "Distributing antidote using pagerank vectors.," *Internet Mathematics*, vol. 6, no. 2, pp. 237–254, 2009.

[11] P. Klepac, O. N. Bjørnstad, C. J. E. Metcalf, and B. T. Grenfell, "Optimizing reactive responses to outbreaks of immunizing infections: Balancing case management and vaccination," *PLoS ONE*, vol. 7, no. 8, p. e41428, 2012.

[12] G. A. Forster and C. A. Gilligan, "Optimizing the control of disease infestations at the landscape scale," *Proc. of the National Academy of Sciences*, vol. 104, no. 12, pp. 4984–4989, 2007.

[13] M. Khouzani, S. Sarkar, and E. Altman, "Optimal control of epidemic evolution," in *Proc. of the 30th Conf. of the IEEE Computer and Communications Societies*, pp. 1683–1691, IEEE, 2011.

[14] C. Borgs, J. Chayes, A. Ganesh, and A. Saberi, "How to distribute antidote to control epidemics," *Random Structures & Algorithms*, vol. 37, no. 2, pp. 204–222, 2010.

[15] P. Van Mieghem, J. Omic, and R. Kooij, "Virus spread in networks," *IEEE/ACM Trans. on Networking*, vol. 17, no. 1, pp. 1–14, 2009.

[16] A. Ganesh, L. Massoulié, and D. Towsley, "The effect of network topology on the spread of epidemics," in *Proc. of the 24th Conf. of the IEEE Comp. and Comm. Societies*, vol. 2, pp. 1455–1466, IEEE, 2005.

[17] J. Díaz, J. Petit, and M. Serna, "A survey of graph layout problems," *ACM Computing Surveys*, vol. 34, pp. 313–356, Sept. 2002.

[18] L. H. Harper, "Optimal assignments of numbers to vertices," *J. of the Society for Industrial and Applied Mathematics*, vol. 12, no. 1, pp. 131–135, 1964.

[19] J. Pantrigo, R. Mart, A. Duarte, and E. Pardo, "Scatter search for the cutwidth mini-

mization problem," *Annals of Operations Research*, vol. 199, no. 1, pp. 285–304, 2012.

[20] E. G. Pardo, N. Mladenović, J. J. Pantrigo, and A. Duarte, "Variable formulation search for the cutwidth minimization problem," *Applied Soft Computing*, vol. 13, no. 5, pp. 2242–2252, 2013.

[21] M. Juvan and B. Mohar, "Optimal linear labelings and eigenvalues of graphs," *Discrete Applied Mathematics*, vol. 36, pp. 153–168, Apr. 1992.

[22] E. Rodriguez-Tello, J.-K. Hao, and J. Torres-Jimenez, "An effective two-stage simulated annealing algorithm for the minimum linear arrangement problem," *Computers & Operations Research*, vol. 35, no. 10, pp. 3331–3346, 2008.

[23] S. Arora, E. Hazan, and S. Kale, "Fast algorithms for approximate semidefinite programming using the multiplicative weights update method," in *Proc. of the 46th IEEE Symp. on Foundations of Computer Science*, pp. 339–348, Oct 2005.

[24] M. Newman, *Networks: An Introduction*. New York, NY, USA: Oxford University Press, 2010.

[25] M. Penrose, *Random geometric graphs*, vol. 5. Oxford University Press, 2003.

[26] J. J. McAuley and J. Leskovec, "Learning to discover social circles in ego networks.," in *NIPS*, pp. 548–556, 2012.

[27] W. A. Massey, "Stochastic orderings for markov processes on partially ordered spaces," *Mathematics of Operations Research*, vol. 12, no. 2, pp. 350–367, 1987.

APPENDIX

Mathematical arguments

In order to simplify the demonstrations, and without loss of generality, we will reorder the nodes of $\mathcal{G}$ according to the priority-order ℓ, and thus consider that $\ell(v_i) = i$.

Notations. First, let $X \in \{0,1\}^N$ be a state vector of size N(i.e. describing the state of the network during the diffusion process), $\mathbb{0}$ and $\mathbb{1}$ vectors of size N that are all-zeros and all-ones, respectively, and $\bar{X} = \mathbb{1} - X$. For $s \subset \{1, ..., N\}$ a subset of nodes, let also $\mathbb{1}_s = (\mathbb{1}_{\{i \in s\}})_i$ be the indicator vector with ones for nodes in the set s. Then, we define τ_X as the extinction time of the contagion starting from the state X, i.e. the time needed for the Markov process to reach its absorbent state $X(t{=}\tau_X) = \mathbb{0}$ when $X(t{=}0) = X$. We also denote the number of contagious nodes in network state X as $N_I(X) = \mathbb{1}^\top X$, while $E_{I-S}(X) = X^\top A\bar{X}$ as the number of edges from a contagious to a healthy node, which edges we also refer to as *contagious edges*.

The proof of Theorem 1 relies on the following lemmas.

Lemma 1. *Under a priority plan, $X \mapsto \mathbb{E}[\tau_X]$ is monotonically increasing with respect to the natural partial order on $\{0,1\}^N$ (i.e. $X \leq Y$ if $X_i \leq Y_i$ for all i).*

Proof. Let $X, Y \in \{0,1\}^N$ be two initial states of the network such that $\forall i \in \{1, ..., N\}, X_i \leq Y_i$. Let also $X(t)$ be the state vector of a contagion initially in state X, i.e. $X(0) = X$, and $Y(t)$ be the state vector of a contagion initially in state Y, i.e. $Y(0) = Y$. This lemma relies on the stochastic domination of $X(t)$ by $Y(t)$. This domination is due to the fact that, under the Markov Process defined by Eq. 1 , and when the control strategy is a priority planning, infection rates are increasing according to the natural partial order on $\{0,1\}^N$, while recovery rates are decreasing. Thus, the initial inequality $X \leq Y$ will, in probability, grow during the contagion process.

The correct proof of this intuition relies on the strong monotonicity of the Markov Process, which we now prove. For all $X \leq Y$, the infection rate of each healthy node at state X is lower than its infection rate at state Y, since $\sum_i A_{ji}X_i \leq \sum_i A_{ji}Y_i$. Also, the recovery rate of each infected node at state X is higher than its recovery rate at state Y: if node i is both infected in state X and Y, and is treated in state Y, then $i = \min\{j \in \{1, ..., N\} : Y_j = 1\}$ is the first infected node, and since the set of infected nodes of state X is included in the set of infected nodes at state Y, we also have $i = \min\{j \in \{1, ..., N\} : X_j = 1\}$ and i is also treated at state X. We can thus apply Theorem 5.4 of [27], and $X(t)$ is a strongly monotone Markov process.

Let $X \leq Y$. If $X(t)$, $Y(t)$ are epidemic processes such that $X(0) = X$ and $Y(0) = Y$, then the strong monotonicity of the Markov process implies that $\forall t \geq 0, \mathbb{P}(\sum_i X_i(t) = 0) \geq \mathbb{P}(\sum_i Y_i(t) = 0)$, which may be rewritten as follows: $\mathbb{P}(\tau_X \leq t) \geq \mathbb{P}(\tau_Y \leq t)$. This means that τ_Y dominates τ_X and thus $\mathbb{E}[\tau_X] \leq \mathbb{E}[\tau_Y]$. □

Lemma 2. *Assume $b_{tot} = 1$ and let X_n^j the worst state vector after j additional infections from the n^{th} state of the planned strategy X_n:*

$$X_n^j = \underset{\substack{\mathbb{1}_s \text{ s.t.} \\ \{n,...,N\} \subset s, \\ |s \cup \{1,...,n-1\}| = j}}{argmax} \mathbb{E}[\tau_{\mathbb{1}_s}]. \tag{7}$$

Then the following bound for the expected extinction time under the priority planning and starting from a total infection holds:
$\forall K \geq 1$ *and* $\rho > \beta \left[\sum_{n=1}^{N} \prod_{j=0}^{K} E_{I-S}(X_n^j)\right]^{\frac{1}{K+1}} - \delta$,

$$\mathbb{E}[\tau_{\mathbb{1}}] \ \leq \ \frac{\sum_{k=0}^{K} f(k)}{(\rho+\delta)(1-f(K+1))}, \tag{8}$$

where

$$f(k) = \sum_{n=1}^{N} \left(\frac{\beta}{\rho+\delta}\right)^k \prod_{j=0}^{k-1} E_{I-S}(X_n^j). \tag{9}$$

Proof. For every state vector X, we have:

$$\mathbb{E}[\tau_X] = \mathbb{E}[t_1 + \tau_{X(t_1)}], \tag{10}$$

where t_1 is the time of the first change in the state vector. Three types of events can happen: either a node recovers by itself (at a rate δ), or a node is healed by a resource (at a rate $\rho + \delta$), or a node is infected (at a rate β). The number of infected nodes is $N_I(X)$. The first contagious node, denoted as $i_X = \min\{j \in \{1,...,N\} : X_j = 1\}$, receives a resource, and the number of nodes that can be infected is $E_{I-S}(X)$. Thus,

$$\begin{aligned}
\mathbb{E}[\tau_X] = \ & \tfrac{1}{\delta N_I(X)+\rho+\beta E_{I-S}(X)} \\
& + \tfrac{\delta(N_I(X)-1)}{\delta N_I(X)+\rho+\beta E_{I-S}(X)} \mathbb{E}[\tau_{X(t_1)} | \text{self-recovery of a node at } t_1] \\
& + \tfrac{\rho+\delta}{\delta N_I(X)+\rho+\beta E_{I-S}(X)} \mathbb{E}[\tau_{X-\mathbb{1}_{i_X}}] \\
& + \tfrac{\beta E_{I-S}(X)}{\delta N_I(X)+\rho+\beta E_{I-S}(X)} \mathbb{E}[\tau_{X(t_1)} | \text{infection at } t_1].
\end{aligned} \tag{11}$$

Using Lemma 1, we get that $\mathbb{E}[\tau_{X(t_1)} | \text{self-recovery of a node at } t_1] \leq \mathbb{E}[\tau_X]$, which leads to:

$$(\delta+\rho+\beta E_{I-S}(X))\mathbb{E}[\tau_X] \leq 1+(\rho+\delta)\mathbb{E}[\tau_{X-\mathbb{1}_{i_X}}]+\beta E_{I-S}(X)\mathbb{E}[\tau_{X(t_1)} | \text{infection at } t_1]. \tag{12}$$

Let $u_n^j = \mathbb{E}[\tau_{X_n^j}]$. For all $j \geq 1$, by definition of X_n^{j-1} and X_n^{j+1}, we have $\mathbb{E}[\tau_{X_n^j - \mathbb{1}_{i_{X_n^j}}}] \leq u_n^{j-1}$ and $\mathbb{E}[\tau_{X_n^j(t_1)} | \text{infection at } t_1] \leq u_n^{j+1}$. This comes from the

fact that, as the order is static, the j infected nodes that are among $\{1, ..., n-1\}$ will receive a resource first. We thus get the following recurrence equation for the u_n^j:

$$(\delta + \rho)(u_n^j - u_n^{j-1}) \leq 1 + \beta E_{I-S}(X_n^j)(u_n^{j+1} - u_n^j), \tag{13}$$

which can be iterated:

$$\begin{aligned}(\rho + \delta)(u_n^0 - u_{n+1}^0) &\leq \textstyle\sum_{k=0}^{K-1} (\frac{\beta}{\rho+\delta})^k \prod_{j=0}^{k-1} E_{I-S}(X_n^j) \\ &\quad + (\rho + \delta)(u_n^{K+1} - u_n^K)(\tfrac{\beta}{\rho+\delta})^K \textstyle\prod_{j=0}^{K} E_{I-S}(X_n^j) \\ &\leq \textstyle\sum_{k=0}^{K-1} (\frac{\beta}{\rho+\delta})^k \prod_{j=0}^{k-1} E_{I-S}(X_n^j) \\ &\quad + (\rho + \delta)u_1^0(\tfrac{\beta}{\rho+\delta})^K \textstyle\prod_{j=0}^{K} E_{I-S}(X_n^j),\end{aligned} \tag{14}$$

since $u_n^{K+1} \leq u_1^0$ using Lemma 1 and $u_1^0 = \mathbb{E}[\tau_{\mathbb{1}}]$.

We can now derive the final formula by summing over n and using the definition $f(k) = \sum_{n=1}^{N} (\frac{\beta}{\rho+\delta})^k \prod_{j=0}^{k-1} E_{I-S}(X_n^j)$:

$$(\rho + \delta)(1 - f(K+1))\mathbb{E}[\tau_{\mathbb{1}}] \leq \sum_{k=0}^{K-1} f(k). \tag{15}$$

□

Lemma 3. *For $n \in \{1, ..., N\}$ and $j \in \{0, ..., n-1\}$,*

$$E_{I-S}(X_n^j) \;\leq\; E_{I-S}(X_n) + j d_{max}, \tag{16}$$

where $d_{max} = \max_i \sum_j A_{ij}$ is the highest degree of the network.

Proof. The contagious nodes of X_n^j are the contagious nodes of X_n and exactly j additional nodes. Since a node can have at most d_{max} neighbors, then each of the j additional nodes can add at most d_{max} edges to the set of contagious edges of the network. □

Lemma 4. *Let $a \geq 0$ and ξ be the (unique) positive solution to $\xi - \ln(1+\xi) = a$. The following inequality holds:*

$$\xi \leq a + 2\sqrt{a}. \tag{17}$$

Proof. $x - \ln(1+x)$ is convex, thus always above its tangent line:
$\forall x_0 > 0$,

$$a = \xi - \ln(1+\xi) \geq (x_0 - \ln(1+x_0)) + \frac{x_0}{1+x_0}(\xi - x_0), \tag{18}$$

and thus,

$$\begin{aligned}\xi &\leq \tfrac{1+x_0}{x_0}(a + \ln(1+x_0)) - 1 \\ &\leq \tfrac{1+x_0}{x_0} a + x_0.\end{aligned} \tag{19}$$

The final result is obtained by setting $x_0 = \sqrt{a}$. □

We can now prove the Theorem 1 using the above lemmas.

Proof of Theorem 1. Using Lemma 2 and Lemma 3, we obtain a bound on the extinction time depending on $\mathcal{C}_{max} = \max_n E_{I-S}(X_n)$,
$\forall K \geq 1$ and $\rho > \beta \left[N \prod_{j=0}^{K} (\mathcal{C}_{max} + j d_{max}) \right]^{\frac{1}{K+1}} - \delta$,

$$\mathbb{E}[\tau_{\mathbb{1}}] \;\leq\; \frac{\sum_{k=0}^{K} f(k)}{(\rho+\delta)(1-f(K+1))} \tag{20}$$

where

$$f(k) = N(\frac{\beta}{\rho+\delta})^k \prod_{j=0}^{k-1} (\mathcal{C}_{max} + j d_{max}), \tag{21}$$

using the approximation $\sum_{n=1}^{N} \prod_{j=0}^{k-1} (E_{I-S}(X_n) + j d_{max}) \leq N \prod_{j=0}^{k-1} (\mathcal{C}_{max} + j d_{max})$.

Finally, we need to select a proper value for K and derive the final result. Let ξ be the unique solution of $\xi - \ln(1+\xi) = \frac{d_{max} \ln N}{\mathcal{C}_{max}}$ and $K^* = \lfloor \frac{\mathcal{C}_{max}}{d_{max}} \xi \rfloor$. Using the particular value of K^*,

$$\begin{aligned}
\sum_{j=0}^{K^*} \ln(1 + j \tfrac{d_{max}}{\mathcal{C}_{max}}) &\leq \int_0^{K^*+1} \ln(1 + x \tfrac{d_{max}}{\mathcal{C}_{max}}) dx \\
&= (K^* + 1 + \tfrac{\mathcal{C}_{max}}{d_{max}}) \ln(1 + (K^*+1) \tfrac{d_{max}}{\mathcal{C}_{max}}) - (K^*+1) \\
&\leq (K^*+1) \ln(1 + (K^*+1) \tfrac{d_{max}}{\mathcal{C}_{max}}) + \tfrac{\mathcal{C}_{max}}{d_{max}} (\ln(1+\xi) - \xi) \\
&= (K^*+1) \ln(1 + (K^*+1) \tfrac{d_{max}}{\mathcal{C}_{max}}) - \ln(N),
\end{aligned} \tag{22}$$

where the second inequality is due to $\frac{\mathcal{C}_{max}}{d_{max}}(K^*+1) \geq \xi$ and the monotonic decrease of $x \mapsto \ln(1+x) - x$ for $x \geq 0$.

From Eq. 22, we derive that $f(K^*+1) \leq \left[\frac{\beta}{\rho+\delta} (\mathcal{C}_{max} + (K^*+1) d_{max}) \right]^{K^*+1}$. We thus have:
For $\rho > \beta(\mathcal{C}_{max} + (K^*+1) d_{max}) - \delta$,

$$\begin{aligned}
\mathbb{E}[\tau_{\mathbb{1}}] \;&\leq \tfrac{\sum_{k=0}^{K^*} f(k)}{(\rho+\delta)(1-f(K^*+1))} \\
&\leq \tfrac{N \sum_{k=0}^{K^*} \left[\frac{\beta}{\rho+\delta} (\mathcal{C}_{max} + (K^*+1) d_{max}) \right]^k}{(\rho+\delta)(1-f(K^*+1))} \\
&\leq \tfrac{N}{(\rho+\delta)(1 - \frac{\beta}{\rho+\delta}(\mathcal{C}_{max} + (K^*+1) d_{max}))} \cdot \tfrac{1 - \left[\frac{\beta}{\rho+\delta} (\mathcal{C}_{max} + (K^*+1) d_{max}) \right]^{K^*+1}}{1 - f(K^*+1)} \\
&\leq \tfrac{N}{\rho + \delta - \beta(\mathcal{C}_{max} + (K^*+1) d_{max})}.
\end{aligned} \tag{23}$$

Finally, using Lemma 4, $d_{max} K^* \leq \mathcal{C}_{max} \xi \leq \mathcal{C}_{max} \left[\frac{d_{max} \ln N}{\mathcal{C}_{max}} + 2 \sqrt{\frac{d_{max} \ln N}{\mathcal{C}_{max}}} \right]$, which proves the desired bound. □